

AD-A037 846

DEFENCE RESEARCH ESTABLISHMENT VALCARTIER (QUEBEC)
USE OF SPATIAL FILTERING TO MATCH WIDE DYNAMIC RANGE GRAYSCALE --ETC(U)
JAN 77 J F BOULTER

F/G 20/6

UNCLASSIFIED

DREV-R-4074/77


NL

1 OF 1
AD
A037846



END

DATE
FILMED
4-77



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

UNCLASSIFIED

CRDV RAPPORT 4074 77
DOSSIER: 3621J-001
JANVIER 1977

DREV REPORT 4074 77
FILE: 3621J-001
JANUARY 1977

3
B.S.

USE OF SPATIAL FILTERING TO MATCH
WIDE DYNAMIC RANGE GRAYSCALE IMAGERY TO
A LOWER RESOLUTION DISPLAY

J.F. Boulter

DDC
RECEIVED
APR 7 1977
RECEIVED
A

DDC FILE COPY

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

Centre de Recherches pour la Défense
Defence Research Establishment
Valcartier, Québec

BUREAU - RECHERCHE ET DEVELOPEMENT
MINISTRE DE LA DEFENSE NATIONALE
CANADA

NON CLASSIFIÉ

RESEARCH AND DEVELOPMENT BRANCH
DEPARTMENT OF NATIONAL DEFENCE
CANADA

CRDV R-4074/77
DOSSIER: 3621J-001

UNCLASSIFIED

14
DREV-R-4074/77
FILE: 3621J-001

②
USE OF SPATIAL FILTERING TO MATCH
WIDE DYNAMIC RANGE GRAYSCALE IMAGERY TO A LOWER RESOLUTION DISPLAY

by

⑩
J. F. Boulter

⑪ Jan 77

⑫ 51 p.

CENTRE DE RECHERCHES POUR LA DEFENSE

DEFENCE RESEARCH ESTABLISHMENT

VALCARTIER

Tel. (418) 844-4271

Québec, Canada

January/janvier 1977

404945

NON CLASSIFIE

13

UNCLASSIFIED

i

RESUME

On examine l'aptitude du filtrage spatial de type non linéaire à réduire la perte d'information quand un système d'affichage possède une plage de résolution dynamique inférieure à celle des niveaux de gris contenus dans l'image qu'on veut visionner. On obtient la compression de la plage dynamique par la réduction des basses fréquences et l'amplification des hautes fréquences en utilisant le filtrage homomorphique. Ce filtrage présuppose que l'image est multipliée par un arrière-plan à basse fréquence. Cette technique, décrite par Oppenheim pour les images formées par la réflexion de la lumière visible, est appliquée, par extension, aux images formées par l'émission de radiation IR et la transmission de rayons X et γ . A l'aide d'exemples, on illustre la compression de la plage dynamique sur ces trois types d'images en utilisant une filtration homomorphique et une procédure d'affichage sans paramètres libres à optimiser en fonction d'une image d'entrée. (NC)

ABSTRACT

We investigate the ability of nonlinear spatial filtering to reduce the amount of information that is lost when a display device has a lower dynamic range than the grayscale imagery to be displayed on it. Dynamic range compression can be achieved by attenuating an assumed low-frequency multiplicative background and enhancing higher frequencies using homomorphic filtering. This technique, originally described by Oppenheim for images formed by reflection of visible light, is extended to images formed by emission of IR radiation and transmission of X-ray and γ -ray radiation. Examples are given which show dynamic range compression of the three types of images achieved by means of a homomorphic filtering and display procedure having no free parameters which must be optimized to suit the particular input image. (U)

UNCLASSIFIED

ii

TABLE OF CONTENTS

RESUME/ABSTRACT.....	i
1.0 INTRODUCTION.....	1
2.0 FILTERING PROCEDURE.....	3
2.1 General Approach.....	3
2.2 Implementation.....	5
2.3 Display Calibration.....	8
3.0 EXAMPLES.....	9
3.1 IRLS Imagery.....	9
3.2 Transmission Radiography.....	17
3.3 Visible-Light Imagery.....	19
4.0 CONCLUSION.....	28
ACKNOWLEDGEMENTS.....	29
REFERENCES.....	30
FIGURES 1 to 7.....	
APPENDIX A Comparison of Homomorphic and Linear Filtering.....	32
APPENDIX B Homomorphic Filtering of IRLS Imagery.....	42
FIGURES A1 to A3 and B1 to B2.....	

UNCLASSIFIED

1

1.0 INTRODUCTION

The dynamic range of imaging sensors such as photographic emulsions, infrared linescan (IRLS) systems, vidicons, etc. can be 30 dB or more. Often, however, a display device or transmission channel with a dynamic range considerably less than that of the imaging sensor is employed, e.g., cathode-ray tubes (CRTs) typically have a dynamic range of less than 15 dB. When imagery from a wide-dynamic-range sensor is to be displayed on a CRT, some form of processing is therefore required to minimize the loss of information.

The usual approach to matching imagery to a display with an insufficient dynamic range is to use a grayscale transformation [1-4]. A wide dynamic range can be compressed by displaying some function of the gray level (such as its logarithm or square root) but this reduces the contrast and signal-to-noise ratio of small-amplitude changes. Level slicing or the use of a look-up table can enhance selected portions of an image but requires some form of operator interaction to achieve the optimum desired result. Automatic techniques such as histogram equilization, contouring, false coloring, etc. have also had limited success in displaying a wide dynamic range and preserving low-contrast details. Displays which dynamically change with time represent an interesting approach but have not yet been well investigated.

Such grayscale processing has the advantage of being relatively easy to implement and can presently be performed digitally at standard

UNCLASSIFIED

2

video rates. However, it is likely that advances in digital signal processing will soon permit more complicated operations, such as spatial filtering, to be performed in real time on high-resolution imagery.

Often the information of interest in a wide-dynamic-range image is contained in the form of edge or high-frequency structure rather than in slow changes in the average brightness level. Here efficient dynamic range compression can be achieved by attenuating the low spatial frequencies corresponding to these changes, while retaining the higher frequency components.

The present report examines the ability of nonlinear high-pass spatial filtering to permit wide-dynamic-range imagery to be displayed on a CRT with a minimum loss of useful information. The processing investigated here, which applies when the low-frequency background and high-frequency structure of interest in the image are multiplied, is an example of homomorphic filtering [5]. In its general form homomorphic filtering is not limited to processing images which are formed as the product of individual parts and can be extended to perform dynamic range compression or enhancement of types of imagery other than those considered here.

After a description of the filtering technique and experimental procedure, results are given for imagery formed by reflection, emission and absorption of radiation (in particular visible-light, IRLS and transmission-radiograph imagery). In each case the filtering significantly increased the amount of information that could be perceived on a CRT display without introducing undesirable artifacts

UNCLASSIFIED

3

such as ringing, enhancement of noise or excessive distortion of the overall appearance of the image. The present filtering procedure has no free parameters which must be adjusted to suit the particular input image, i.e. all images are processed in an identical manner.

This work was performed at DREV during 1975-76 under PCN 15B22 Image Processing and PCN 21J01 Guidance and Control Technology.

2.0 FILTERING PROCEDURE

2.1 General Approach

The low- and high-frequency parts of an image are seldom additive so, in general, simple high-pass linear filtering cannot be used to attenuate the low spatial frequencies. An explanation of why this is so is given in Appendix A. Oppenheim et al [5] have described a procedure which, under certain conditions, can be used to simplify the required non-linear filtering. The technique (homomorphic filtering) applies for signals formed by combining individual components according to any type of operator that can be made to obey superposition. The approach is to transform the signal so that the components to be individually filtered become additive, apply the required linear filter and take the inverse transform. An example of such an operator is multiplication, the transform and inverse transforms being the logarithmic and exponential functions respectively.

UNCLASSIFIED

4

The use of homomorphic high-pass filtering to perform dynamic range compression, originally applied by Oppenheim [5] to images formed by reflection of visible light, is extended here to images formed by emission of IR radiation and absorption of X-ray and γ -ray radiation. In each case high-frequency detail of interest is assumed to be present along with a slowly varying background which covers a wide dynamic range. The wide dynamic range required to display the background causes the higher frequency information to appear with low contrast; the filtering attempts to attenuate this background with minimum loss or distortion of useful information.

To determine a single filtering procedure suitable for the present three types of imagery, we make the assumption that the low-frequency background to be removed is multiplied by, rather than added to, the high-frequency structure. This is based primarily on the simplified models for the formation of each type of image described in Sec. 3. Secondly, the filtering required to remove a low-frequency multiplicative background always produces a non-negative result, as is required on physical grounds, while the filtering to remove an additive background does not (Appendix A). Dynamic range compression is then achieved for the present imagery by calculating the logarithm, attenuating the low spatial frequencies in an appropriate fashion and exponentiating.

Several low-frequency cutoff characteristics, including sharp cutoff, gaussian, $w^{1/4}$, w and w^2 ($w^2 = u^2 + v^2$ where u and v are spatial frequencies) were investigated experimentally. Images with the most resolved detail and least apparent distortion or noise enhancement were

UNCLASSIFIED

5

obtained by attenuating the low frequencies according to $w^{1/2}$. In most images there was considerable overlap of the power spectra of the low- and high-frequency components so that a smooth and gradual low-frequency attenuation, which minimized overshoot or ringing effects, was preferable to a sharper cutoff.

It has been suggested that similar multiplicative processing occurs in the human visual system. If a high-pass filter is assumed to act on the perceived image after it has been logarithmically sensed by the eye, this can qualitatively explain some of the characteristics and illusions present in human vision [6]. In this context it is interesting to note that when the present nonlinear $w^{1/2}$ filter, which was selected on a subjective basis, was applied to appropriate patterns, effects similar to Mach banding and simultaneous contrast [6], which are characteristics of human vision, were produced. The present processing could then be considered as an attempt to duplicate that occurring in the human visual system. However, it is performed before the imagery has been degraded by display on a device with insufficient dynamic range.

2.2 Implementation

Each image was originally in the form of a positive or negative film transparency. Over its normal range the density $D(x,y)$ of such a transparency is related to the exposure $E(x,y)$ according to:

$$D(x,y) = \gamma \log E(x,y) \quad (1)$$

UNCLASSIFIED

6

where the parameter γ depends on the properties of the film and its development. The transmission through the film $T(x,y)$ is obtained by exponentiating the density:

$$T(x,y) = 10^{D(x,y)} = E^{\gamma}(x,y) \quad (2)$$

The film therefore has a power-law response with the transmission obtained by raising the input brightness to the power γ . In the present processing the transmission image $T(x,y)$ rather than the true exposure $E(x,y)$ is taken as the "original" image.

Using an Optronics P-1700 microdensitometer, each film transparency was digitized with a 25-, 50- or 100- μm raster into 256 uniformly spaced density levels covering a k-decade density range. The value $k=2$, corresponding to the density range 0-2D, was used unless otherwise stated. A 512- by 512-element region of the digitized image was then selected for processing. Where necessary the polarity of the images was reversed after the logarithmic digitization so that high exposure levels in the original scene always appeared as large digitized values.

The digitized image was divided by 255 to give, for negative film, a normalized density image $f_d(x,y)$ related to the film transmission by:

$$f_d(x,y) = \frac{-1}{k} \log T(x,y) \quad (3)$$

UNCLASSIFIED

7

A similarly normalized transmission image $f_t(x,y)$ was obtained by exponentiating the density image:

$$f_t(x,y) = 10^{kf_d(x,y)} \quad (4)$$

The filtering applied to $f_d(x,y)$ was implemented in the frequency domain using the two-dimensional fast Fourier transform (FFT) [7]. Edge effects were minimized by matching the left and right, and upper and lower edges of $f_d(x,y)$ by linear interpolation [8] before transforming to the frequency domain. The matching was performed over a region 5 or 10 elements wide bordering the image. The two-dimensional FFT of the edge-matched image was multiplied by the filter function:

$$Q(u,v) = 2\sqrt{u^2+v^2} \quad (5)$$

where u and v are in units of the Nyquist frequencies (i.e., $0 \leq u,v \leq 1$). The inverse FFT was calculated and the result exponentiated according to Eq. 4 to give the processed transmission image. Photographs of the original and processed transmission images were then made according to the procedure described in Sec. 2.3.

The processing is equivalent to the use of a film with a frequency-dependant γ which approaches a value of two at high frequencies and 0 at low frequencies. The high frequencies in the processed transmission image are therefore enhanced (the amplitudes of the highest frequencies are squared) while low spatial frequencies are attenuated (the DC component is set to zero).

2.3 Display Calibration

The improvement in subjective quality of the image as recorded on the film should represent the effect of the spatial filtering rather than optimization of display parameters to suit the particular input image.

Before displaying each original or processed image, all gray levels differing from the mean by more than 3σ (where σ is the standard deviation of the grayscale distribution) were clipped. The contrast was then expanded so that the minimum gray level remaining in the image was displayed as full black and the maximum gray level as full white. Such clipping and normalization usually had a negligible effect when applied directly to the original image. However it significantly improved the contrast of the processed image by limiting the amplitudes of high-contrast points which were enhanced by the filtering. (Retention of these points usually caused the grayscale range of interest in the image to be displayed with lower contrast.)

All images were then displayed on a Tektronix 602 x-y-z CRT (with approximately 300- by 300-element spatial resolution) and photographed using Polaroid type 47 film. The camera exposure time and aperture, and CRT brightness and contrast were adjusted to obtain a photograph in which the black band of a 16-level linear grayscale chart was barely exposed while the white band was near film saturation. When calibrated in this manner the black band was visible in a darkened room while the white band was approximately 75% of the CRT's saturation brightness.

This exposure calibration and grayscale normalization ensured that most of the dynamic range of the CRT was employed and recorded on the Polaroid film for both the original and processed images. The grayscale quality of the photographs was limited by noise and nonuniformity in the CRT phosphor rather than by the film.

3.0 EXAMPLES

3.1 IRLS Imagery

To a first approximation an IR detector is sensitive to the product of surface emissivity $e(x,y)$ and a term dependant on the absolute temperature $T(x,y)$:

$$a(x,y) = c e(x,y) f[T(x,y)] \quad (6)$$

Variations in both emissivity and temperature therefore contribute to the radiation contrast, and in some cases the two effects can be of equal magnitude [9].


Large variations in the local average value of $a(x,y)$ often occur in an IR image, for example, at a land-sea boundary. The large dynamic range required to display these changes can cause higher frequency structure to appear with low contrast. To attenuate the slowly varying changes in temperature, while retaining higher frequency structure in emissivity, the two parts are made additive by taking the logarithm:

$$\log a(x,y) = \log c + \log e(x,y) + \log f[T(x,y)] \quad (7)$$

a high-pass filter is applied and the result is exponentiated.

UNCLASSIFIED

10




(a)

FIGURE 1 Compression of the dynamic range of an IRLS image by spatial filtering. The CRT displays of the original and processed images are shown in (a) and (b) respectively. A histogram giving the distribution of gray levels is shown below each image.

UNCLASSIFIED

11



(b)

FIGURE 1 Compression of the dynamic range of an IRLS image by spatial filtering. The CRT displays of the original and processed images are shown in (a) and (b) respectively. A histogram giving the distribution of gray levels is shown below each image.

UNCLASSIFIED

12

Two IRLS images obtained using a Reconofax camera (8-14 μm band) are shown in Figs. 1(a) and 2(a). Because the IR detector signal is AC coupled and has an arbitrary DC bias level added to it, certain conditions must be satisfied if the logarithmic function is to make the emission and temperature parts of an IRLS image additive (Appendix B). However, as discussed in Appendix B, the filtering required to remove a multiplicative low-frequency background from IRLS imagery is preferable to that required to remove an additive background.

The result of processing the first image is shown in Fig. 1(b). A histogram, giving the distribution of gray levels, is included below each image. The grayscale distribution of the original image contains a peak corresponding mainly to the darker lower portion of the image, along with a relatively flat background extending to higher brightness levels. As is evident in the histogram of Fig. 1(b), the processing has distributed the gray levels more evenly (almost log-normally in transmission or normally in density) around a central value, yielding an image with improved contrast and definition. For example, paths near the buildings and foliage in the lower dark region of the image have been strongly enhanced.

The second IRLS image, Fig. 2(a), contains variations in brightness which occur over a much smaller scale than in Fig. 1(a). The filtering, however, achieved the same overall result with high-frequency detail appearing with improved contrast, light areas being decreased in brightness and dark areas increased in brightness (Fig. 2(b)).

UNCLASSIFIED

13

The enhancement of high frequencies and overall sharpening evident in Figs. 1(b) and 2(b) was achieved without excessive enhancement of noise or introduction of features not evident in the original film transparency. For example, no ringing or overshoot occurred even at the sharp highway boundaries present in both images.

As can be seen in the histogram of the IRLS image of Fig. 1(a), a significant fraction of all image elements have grayscale values within a relatively narrow peak region. One might therefore expect that this image could be enhanced effectively using standard grayscale transform techniques. Three such techniques were applied to this image for comparison with the spatial filtering.

For the first grayscale processing, shown in Fig. 3(b), the grayscale from 0 to 0.25 was stretched to cover the full range from black to white. All levels greater than 0.25 were saturated full white. The second processed image, Fig. 3(c), employed a four-cycle repeating grayscale implemented by setting the two most significant bits of each image element to zero and multiplying by four. In the third, Fig. 3(d), the grayscale was compressed using a two-decade logarithmic display (resulting in an image which represents film density rather than transmission). The original image is repeated in Fig. 3(a), and the result of applying the processing to a linear grayscale chart is shown below each respective image to illustrate the grayscale transfer function.

UNCLASSIFIED

14



(a)

FIGURE 2 Compression of the dynamic range of an IRLS image by spatial filtering. The CRT displays of the original and processed images are shown in (a) and (b) respectively.

UNCLASSIFIED

15



(b)

FIGURE 2 Compression of the dynamic range of an IRLS image by spatial filtering. The CRT displays of the original and processed images are shown in (a) and (b) respectively.

UNCLASSIFIED

16




(a)



(b)



(c)



(d)

FIGURE 3 Dynamic range compression of an IRLS image using 3 standard grayscale transforms. The original image (a) is displayed in (b) with the grayscale range 0 to 0.25 stretched, in (c) using a 4-cycle repeating grayscale and in (d) by compressing the dynamic range logarithmically. The result of applying the transforms to a linear grayscale chart is shown below each respective image.

While effectively enhancing the lower dark region of the original image, the grayscale stretching removed much useful information in other brighter areas. The four-cycle repeating grayscale discarded less information but yielded a display in which certain patterns, such as the buildings above the highway, could no longer be recognized because of the false contours that were introduced. The use of a two-cycle repeating grayscale yielded a result only slightly less ambiguous in the upper half of the image but with the contrast in the lower half reduced by a factor of two. The logarithmic compression retained the full dynamic range of the grayscale but at the expense of a considerable loss in contrast. In comparing Figs. 3(b), (c) and (d) with Fig. 1(b), it is evident that for the present image the three grayscale transforms are much less acceptable than the spatial filtering technique for performing efficient dynamic range compression.

3.2 Transmission Radiography

In transmission radiography an image is formed by recording, usually on film, radiation which has been attenuated by passing through an object. The structure of principal interest in such imagery (e.g., cracks or defects) may correspond to only very small changes in attenuation while the film is required to record the wide dynamic range caused by large variations in the thickness of the object. Attenuation of the slow thickness variations therefore allows higher frequency structure in the radiograph to be displayed with improved contrast.

Assume that the object to be radiographed is cut into N thin slices perpendicular to the direction of the incident radiation and that

the absorption coefficient integrated over the thickness of the i th slice yields the function $g_i(x,y)$. Neglecting scattering and other second-order effects, a ray entering the object at (x,y) with initial intensity I_0 will be attenuated to the value:

$$I(x,y) = I_0 e^{-g_1(x,y)} e^{-g_2(x,y)} \dots e^{-g_N(x,y)} \quad (8)$$

A single slice may contain low-frequency structure (if it corresponds to a uniform interior region or a boundary region with smooth edges) or higher frequency structure (if it corresponds to non uniformities in the interior or to rough surfaces). The final image $I(x,y)$ is formed by multiplying the attenuation $\exp[-g_i(x,y)]$ contributed by each of the N slices. Taking the logarithm of Eq. 8 makes the individual contributions additive:

$$\ln I(x,y) = \ln I_0 - g_1(x,y) - g_2(x,y) - \dots - g_N(x,y) \quad (9)$$

and therefore of a form in which a linear high-pass filter can be used to attenuate the slow thickness variations. Such filtering also applies for other types of images which are formed as a result of exponential absorption of radiation.

A γ -ray radiograph of an artillery shell was digitized over the density range 1D to 4D. The CRT displays of the original and processed images are shown in Figs. 4(a) and (b). Structural information such as the outline of the explosive charge within the shell as well as other high-frequency and edge details not visible in the original image are immediately apparent after processing. A graph of row number 150

(counting upward from the bottom) is shown below each 512- by 512-element image. The graph from the original image shows high-frequency structure superimposed on a slowly varying background component which covers a wide dynamic range. In the processed image the background component has been strongly attenuated allowing the high-frequency structure to appear with greater amplitude.

A second transmission radiograph, a dental X-ray, was similarly processed as shown in Fig. 5. Again, the dynamic range compression has increased the visibility of structure in dark regions (such as blood vessels between the teeth) as well as in bright regions (details of the metal cap in the upper left of the image).

3.3 Visible Light Imagery

A visible-light image $g(x,y)$ can be formed as the product of the reflectivity of the surface $r(x,y)$ and the illumination $b(x,y)$. If the illumination varies slowly over the target then it can be attenuated relative to higher frequency reflectivity information by calculating the logarithm to make the two components additive:

$$\log g(x,y) = \log r(x,y) + \log b(x,y) \quad (10)$$

applying high-pass filtering and then exponentiating the result. Oppenheim et al [5] showed how visible-light imagery could effectively be enhanced by such filtering.

UNCLASSIFIED

20





FIGURE 4 Compression of the dynamic range of a transmission radiograph image by spatial filtering. The CRT displays of the original and processed images are shown in (a) and (b) respectively. A graph of row number 150 (counting from the bottom) is given below each image.

UNCLASSIFIED

21



(b)


FIGURE 4 Compression of the dynamic range of a transmission radiograph image by spatial filtering. The CRT displays of the original and processed images are shown in (a) and (b) respectively. A graph of row number 150 (counting from the bottom) is given below each image.

UNCLASSIFIED

22



(a)



(b)

FIGURE 5 Compression of the dynamic range of a dental X-ray by spatial filtering. The CRT displays of the original and processed images are shown in (a) and (b) respectively.

The visible-light image of an APC vehicle is shown in Fig. 6(a). The result of filtering this image, using the same procedure followed for the IRLS and transmission radiograph imagery, is given in Fig. 6(b). The processed image is considerably sharper than the original with a "flat-lighting" effect that makes details in the shadowed regions more readily visible. For example, the trees and shadowed regions near the wheels have been increased in brightness and contrast while strongly illuminated areas such as the sky and the lower-left portion of the road, have been reduced in brightness.

A second example of the homomorphic filtering of a visible-light image is shown in Fig. 7. The original image (Fig. 7(a)) shows the shadow of a cloud extending across a portion of the Grand Canyon of Arizona. After filtering (Fig. 7(b)) details in both the shadowed and the strongly-illuminated regions have been enhanced. In the original image very little structure is apparent in the shadow of the cloud while after the dynamic range compression features (such as the river) are clearly visible.

UNCLASSIFIED

24



(a)

FIGURE 6 Compression of the dynamic range of a visible-light image by spatial filtering. The CRT displays of the original and processed images are shown in (a) and (b) respectively.

UNCLASSIFIED

25



(b)

FIGURE 6 Compression of the dynamic range of a visible-light image by spatial filtering. The CRT displays of the original and processed images are shown in (a) and (b) respectively.

UNCLASSIFIED

26



(a)

FIGURE 7 Compression of the dynamic range of a visible-light image by spatial filtering. The CRT displays of the original and processed images are shown in (a) and (b) respectively.

UNCLASSIFIED

27



(b)

FIGURE 7 Compression of the dynamic range of a visible-light image by spatial filtering. The CRT displays of the original and processed images are shown in (a) and (b) respectively.

UNCLASSIFIED

28

4.0 CONCLUSION

Compared to grayscale transformation, spatial filtering is a powerful technique for compressing the dynamic range of imagery to suit a lower-dynamic-range transmission channel or display device. In the present report homomorphic high-pass filtering was shown to be capable of effectively attenuating a slowly varying multiplicative background in IRLS, transmission radiograph and visible-light imagery. This permitted higher frequency structure to appear with a better contrast and signal-to-noise ratio on a low dynamic range CRT display.

Each processed image retained the overall appearance of the original but with edges and high-frequency structure considerably sharpened. The enhancement was achieved without introducing undesirable artifacts (such as ringing or overshoot) or significantly increasing the noise apparent in the image. It can be emphasized that each image was processed and displayed in an identical manner. Optimization or adjustment of parameters to suit the particular input image was not required.

The present homomorphic filtering attempted specifically to reduce the amount of information that is lost when a CRT has a lower dynamic range than the image to be displayed on it. However, the ability of the nonlinear processing to sharpen IRLS and transmission radiograph images, possibly making them more suitable for either human or machine analysis, was also demonstrated.

UNCLASSIFIED

29

Extension of homomorphic filtering to perform more general types of image enhancement is possible. In particular the potential of using homomorphic filtering to compress the dynamic range or enhance selected parts of "non-natural" images (i.e. two-dimensional grayscale displays such as voiceprints, sonagrams etc.) has not been well explored. Such images may not necessarily be multiplicative in nature and could require a transform other than the present logarithm-exponentiate operation (such as integrate-differentiate) to make the individual parts additive.

ACKNOWLEDGEMENTS

Useful and interesting discussions with Dr G. Giroux, Mr G. A. Morley and Mr D. S. Galbraith are gratefully acknowledged. The two-dimensional FFT routines were programmed by Mr L. Sévigny.

UNCLASSIFIED

30

REFERENCES

1. Andrews, H. C., "Monochrome Digital Image Enhancement", Appl. Opt., Vol. 15, pp. 495-503, February 1976.
2. Huang, T. S., Schreiber, W. F. and Tretiak, O. J., "Image Processing", Proc. IEEE, Vol. 59, pp. 1586-1609, November 1971.
3. Hummel, R. A., "Histogram Modification Techniques", Computer Graphics and Image Processing, Vol. 4, pp. 209-224, 1975.
4. Andrews, H. C., "Digital Image Restoration: A Survey", IEEE Computer, Vol. 7, pp. 36-45, May 1974.
5. Oppenheim, A. V., Schaffer, R. W. and Stockham, T. G. Jr., "Nonlinear Filtering of Multiplied and Convolved Signals", Proc. IEEE, Vol. 56, pp. 1264-1291, August 1968.
6. Stockham, T. G., Jr., "Image Processing in the Context of a Visual Model", Proc. IEEE, Vol. 60, pp. 828-842, July 1972.
7. Boulter, J. F., "Use of Two-Dimensional Digital Fourier Transforms for Image Processing and Analysis", DREV R-4025/75, July 1975, Unclassified.
8. McDonnell, M. J. and Bates, R. H. T., "Preprocessing of Degraded Images to Augment Existing Restoration Methods", Computer Graphics and Image Processing, Vol. 4, pp. 25-39, 1975.

UNCLASSIFIED

31

9. Hudson, R. D., Jr., "Infrared System Engineering", John Wiley and Sons, New York, 1969.

UNCLASSIFIED

32

APPENDIX A

Comparison of Homomorphic and Linear Filtering

In the present report the dynamic range of imagery is reduced by high-pass spatial filtering. The filtering assumes that the image to be filtered is composed of a part containing mainly low spatial frequencies along with a part containing mainly high spatial frequencies. The filtering attempts to attenuate the low-frequency part, which is assumed to contain little useful information, while retaining the high-frequency part of interest. The single model found most suitable for the three types of imagery which were considered (in particular images formed by reflection, emission and absorption of radiation) is that the low-frequency part is multiplied by, rather than added to, the high-frequency part. This appendix shows why, for the present imagery, the low-frequency part is better attenuated using homomorphic high-pass rather than conventional linear high-pass filtering.

A simple example will illustrate the difference between homomorphic and linear high-pass filtering of multiplied signals. A low-frequency two-dimensional biased sine wave:

$$A = a + \sin w_1 r$$

UNCLASSIFIED

33

is multiplied by a high-frequency two-dimensional biased sine wave:

$$B = b + \sin w_h r$$

where $w_h \gg w_l$, to form the image to be filtered:


$$P = AB = (a + \sin w_l r)(b + \sin w_h r) \quad (A1)$$

The spatial frequencies w_l and w_h and the space variable r are vectors which describe two-dimensional sinusoidal surfaces and, in accord with the present types of images which are formed as the product of positive-only parts, the constants a and b are greater than 1.


An example of such an image is shown in Fig. A1(c). It was formed by multiplying a low-frequency sine wave with $u=v=4$ cycles/picture width (Fig. A1(a)) by a high-frequency sine wave with $u=-v=17$ cycles/picture width (Fig. A1(b)). Both sine waves were biased with a level $a=b=2$ before being multiplied. (The contrast of these and all subsequent photographs is improved by ignoring any DC level and displaying the minimum level in the image as full black and the maximum level as full white.) The image of Fig. A1(c) will be filtered using both linear and homomorphic high-pass filters to remove the low-frequency part A with minimum distortion of the high-frequency part B.

UNCLASSIFIED


34



(a) low-frequency part




(b) high-frequency part



(c) product of low- and high-frequency parts

FIGURE A1 A low-frequency two-dimensional biased sine wave (a) is multiplied by a high-frequency two-dimensional biased sine wave (b) to form the image to be filtered shown in (c).



(a) power spectrum of the image



(b) power spectrum of the logarithm of the image

FIGURE A2 The power spectrum of the image to be filtered (a) contains sum and difference frequencies in addition to the fundamental low- and high-frequency components. The power spectrum of the logarithm of the image to be filtered (b) contains no sum and difference frequencies but higher order harmonics of the fundamental low- and high-frequency components have been introduced by the logarithmic function.

The linear filter is first applied. Expanding Eq. A1 yields sum and difference frequencies in addition to the fundamental low- and high-frequency terms:

$$P = ab + b \sin w_1 r + a \sin w_h r + 1/2 \cos (w_1 - w_h) r - 1/2 \cos (w_1 + w_h) r \quad (A2)$$

The four individual frequency components can be seen in the power spectrum of Fig. A1 (c) which is shown in Fig. A2(a). Considering only the upper half of this power spectrum (the lower half is its conjugate) the dot close to the center of the image represents the fundamental low-frequency component w_1 while the middle dot in the group of three


dots represents the fundamental high-frequency component w_h . The sum and difference frequencies appear on either side of the high-frequency component. (The DC component which should appear at the center of the photograph has been removed and a 5-decade logarithmic grayscale has been used to improve the display of this and subsequent power spectra.)

Linear high-pass filtering can correctly attenuate the DC and fundamental low-frequency component but not the sum and difference frequencies since they are within the high-frequency range of interest that must be retained. This fact (i.e. the presence of the sum and difference frequencies) represents the basic problem in attempting to process multiplied signals using a linear filter. Applying a linear filter to Eq. A2 to eliminate all frequencies at or below w_1 yields:


$$P(\text{linear}) = a \sin w_h r + 1/2 \cos (w_1 - w_h)r - 1/2 \cos (w_1 + w_h)r \quad (\text{A3})$$

The result of linearly filtering Fig. A1(c) in this way is shown in Fig. A3(a). Due to the presence of the sum and difference frequencies (each with an amplitude $1/2a$ of the desired result $\sin w_h r$) the filtered image is quite different from the pure high-frequency component shown in Fig. A1(b) and only slightly improved over the original image Fig. A1(c). For the present case with $a=b=2$, the amplitudes of the sum and difference frequencies are $1/4$ that of the desired high-frequency component.

Next, the homomorphic high-pass filter is applied. This is performed by calculating the logarithm of the image to make the two



(a) linear high-pass filter



(b) homomorphic high-pass filter

FIGURE A3 The image shown in Fig. A1(c) is processed using a linear high-pass filter (a) and a homomorphic high-pass filter (b) to remove the low-frequency part. The linear-filtered image is only slightly improved over the original because of retention of the sum and difference frequencies while the homomorphic-filtered image is indistinguishable from the desired result, Fig. A1(b).

parts additive, applying a linear high-pass filter and then exponentiating. The logarithm of Eq. A1 is:

$$\begin{aligned} \ln P &= \ln A + \ln B \\ &= \ln (a + \sin w_1 r) + \ln (b + \sin w_h r) \end{aligned} \quad (A4)$$

To third order a series expansion for the logarithmic function yields:

$$\begin{aligned} \ln P \approx & [\ln a + \ln b - 1/4a^2 - 1/4b^2] + [1/a + 1/4a^3] \sin w_1 r \\ & + 1/4a^2 \cos 2 w_1 r - 1/12a^3 \sin 3 w_1 r + [1/b + 1/4b^3] \sin w_h r \\ & + 1/4b^2 \cos 2 w_h r - 1/12b^3 \sin 3 w_h r \end{aligned} \quad (A5)$$

The sum and difference frequencies present during the linear filtering case are absent, however higher-order harmonics of w_1 and w_h have been introduced by the non-linear logarithmic function. The power spectrum of the logarithm of Fig. A1(c) is shown in Fig. A2(b). The first- and second-order harmonics of w_1 and w_h (located along the line passing through the fundamental component and the DC origin) are of sufficient amplitude to be visible in the photograph. Ideally we want the high-pass filter to remove w_1 and all its higher harmonics and to retain w_h and its harmonics. Applying such a filter to Eq. A5 and exponentiating yields (to third order in the series expansion for the exponential function):

$$P(\text{ideal homomorphic}) = 1 + 1/b \sin w_h r \quad (\text{A6})$$

Within a gain factor this is equal to the desired high-frequency biased sine wave B. In addition, in contrast to the linear-filtered result given in Eq. A3, the homomorphic filter ensures that the result is always positive as is required on physical grounds.

A high-pass filter cannot eliminate all harmonics of the low-frequency term without also removing the high-frequency terms of interest. The magnitude of the error caused by incomplete removal of the harmonics of w_1 can be estimated by assuming that the high-pass filter only eliminates frequencies at or below the first harmonic of w_1 (i.e. harmonics equal to or higher than $3w_1$ are retained). Application of this filter to Eq. A5 followed by exponentiation yields (to third order):

$$P(\text{homomorphic}) = 1 + 1/b \sin w_h r - 1/12a^3 \sin 3 w_1 r \quad (A7)$$

Thus an undesired term at $3w_1$ with an amplitude $b/12a^3$ times that of the fundamental high-frequency component remains. For the present values ($a=b=2$) the amplitude of the $3w_1$ term is only $1/48$ that of the desired w_h term. The result of filtering the image of Fig. A1(c) using this homomorphic filter is shown in Fig. A3(b). It is evident that the low-frequency part has been effectively removed with little distortion of the high-frequency part.

Neither the linear nor the homomorphic filter can perfectly remove the low-frequency part from the image of Fig. A1(c). Both leave undesired error terms in the final image. For the linear filter the error terms are the sum and difference frequencies caused by the multiplication while for the homomorphic filter they are the harmonics caused by the logarithmic function. In the present example the amplitude of the error present after the homomorphic filtering is smaller than that present after the linear filtering by the factor $6a^2/b$. For the image of Fig. A1(c) with $a=b=2$ the error is 12 times smaller for the homomorphic filter. In addition, as contrasted with the high-frequency nature of the linear-filter error, the homomorphic-filter error is predominantly low frequency and therefore more easily reduced by using an appropriate low-frequency cutoff characteristic.

Practical dynamic range compression requires the use of a filter with a gradual low-frequency cutoff to minimize ringing or overshoot effects in the final image. In Sec. 3 of the present report a filter with a cutoff which varied as the square-root of the spatial frequency was applied to the logarithm of the image. Such a filter

UNCLASSIFIED

40

attenuates the lowest spatial frequencies along with their lower-order harmonics (higher harmonics of a fundamental low-frequency component caused by the logarithmic function decrease rapidly in amplitude and therefore require less attenuation). Since the power spectrum of many "complicated" images increases rapidly at low frequencies only the lowest frequencies and their lower-order harmonics may have to be strongly attenuated to achieve significant dynamic range compression.

In a complicated image where the two parts will both contain more than a single spatial frequency, $1/a$ then represents the amplitude of an arbitrary low-frequency component present in A, and $1/b$ the amplitude of an arbitrary high-frequency component present in B, relative to the amplitude of the DC component. The above error analysis can then be applied to such complicated images if the single low- and high-frequency terms in Eq. A1 are replaced by summations over a range of spatial frequencies each with a given amplitude and phase.

The present report has assumed that the low- and high-frequency parts of the image were independent of one another and had well separated power spectra. This will not be true in general for complicated images. Spectral overlap will most certainly occur and the low- and high-frequency parts can be correlated (e.g. in an infrared image the temperature of a region can depend on its emissivity). The purpose here has been to find a single filtering procedure with no adjustable parameters which could effectively compress the dynamic range of widely differing types of imagery, in particular those formed by

UNCLASSIFIED

41

emission, absorption or reflection of radiation. Assuming that the dynamic range compression is to be accomplished by attenuation of a low-frequency multiplicative background, this appendix has shown why homomorphic filtering is preferable to linear filtering.

UNCLASSIFIED

42

APPENDIX B

Homomorphic Filtering of IRLS Imagery

To form an IRLS image, the output of a scanned IR detector is AC coupled, biased to ensure a positive signal and fed to a lamp, synchronized with the detector scan, which exposes a film. Assume that the film exposure is proportional to the voltage driving the lamp and that the detector output is proportional to the product of a low-frequency "temperature image" A and a high-frequency "emissivity image" B. Further assume that the AC coupling time constant is long enough that it does not remove low-frequencies of interest from the image. For an airborne IRLS system, this means that the equivalent spatial "time constant" in the direction of aircraft motion resulting from the AC coupling is much longer than typical features of interest in the image. (The power spectrum of an IRLS image taken from the same strip of film as Fig. 1(a) displayed no significant loss of low spatial frequencies indicating that this condition on the AC coupling time constant is probably satisfied for Fig. 1(a). This power spectrum is shown in Fig. 15(b) of [7].)

The effect of the AC coupling is then to remove the DC level averaged over a region of the image which is larger than typical features of interest. A region from which an average DC level k_1 has been removed by the AC coupling and to which a bias level k_2 has been added will have a film exposure:

UNCLASSIFIED

43

$$E = AB - k_1 + k_2$$

If the bias level is chosen equal to the average DC level removed by the AC coupling, the film exposure is equal to AB . The previous homomorphic filter for multiplicative images (i.e. logarithm, linear high-pass filter, exponential) will then correctly attenuate the low-frequency temperature image A . However, if the error in the DC level ($k_1 - k_2$) is significant, the present homomorphic filter no longer applies (the logarithmic function does not make the parts A and B additive).


We will not attempt to analyse here what effect an arbitrary DC level in an IRLS image will have on the present homomorphic filtering. Instead two examples will be given which show how a normal linear filter and the present homomorphic filter compare in this situation.

First a DC level with an amplitude of 0.5 was added to the simulated image of Fig. A1(c). This image consists of a low-frequency biased sine wave multiplied by a high-frequency biased sine wave and is described by Eq. A1 with $a=b=2$. The DC level which was added then represents an error equal to half the amplitude of the DC component present in the original image. The results of filtering this image to remove the low-frequency component are shown in Fig. B1. The same linear and homomorphic filters that were employed to obtain Figs A3(a) and (b) were used. While structure related to the undesired high-frequency component is evident in the homomorphic-filtered image (Fig. B1(b)), the result is still superior to that obtained using the linear filter (Fig. B1(a)).


UNCLASSIFIED

44

The second example, given in Fig. B2(a), shows the result of filtering the IRLS image of Fig. 1(a) using a linear filter which has the same low-frequency cutoff characteristic (i.e. square-root of the spatial frequency) as was used for the homomorphic filter described in Sec. 2. For comparison, the homomorphic-filtered image obtained in Sec. 3 is included in Fig. B1(b). The overall contrast of the homomorphic-filtered IRLS image and the degree of dynamic range compression achieved is superior to that of the linear-filtered image. For example, the sharpness of the boundaries of the horizontal highway near the center of the image is considerably better in the homomorphic-filtered image and more details are evident in the lower half.




(a) linear high-pass filter




(b) homomorphic high-pass filter

FIGURE B1 Use of linear and homomorphic high-pass filters to remove a low-frequency multiplicative component from an image which contains an arbitrary DC level



(a) linear high-pass filter



(b) homomorphic high-pass filter

FIGURE B2 Comparison of the abilities of linear and homomorphic high-pass filters to compress the dynamic range of an IRLS image containing an arbitrary DC level

DREV R-4074/77 (UNCLASSIFIED)

Bureau- Recherche et Développement, Ministère de la Défense nationale, Canada.
CRDV, C.P. 880, Courcellette, Qué. G0A 1R0

"Use of Spatial Filtering to Match Wide Dynamic Range Grayscale to a Lower Resolution Display" par J.F. Boulter

On examine l'aptitude du filtrage spatial de type non linéaire à réduire la perte d'information quand un système d'affichage possède une plage de résolution dynamique inférieure à celle des niveaux de gris contenus dans l'image qu'on veut visionner. On obtient la compression de la plage dynamique par la réduction des basses fréquences et l'amplification des hautes fréquences en utilisant le filtrage homomorphe. Ce filtrage présuppose que l'image est multipliée par un arrière-plan à basse fréquence. Cette technique, décrite par Oppenheim pour les images formées par la réflexion de la lumière visible, est appliquée, par extension, aux images formées par l'émission de radiation IR et la transmission de rayons X et γ . A l'aide d'exemples, on illustre la compression de la plage dynamique sur ces trois types d'images en utilisant une filtration homomorphe et une procédure d'affichage sans paramètres libres à optimiser en fonction d'une image d'entrée. (NC)

DREV R-4074/77 (UNCLASSIFIED)

Bureau- Recherche et Développement, Ministère de la Défense nationale, Canada.
CRDV, C.P. 880, Courcellette, Qué. G0A 1R0

"Use of Spatial Filtering to Match Wide Dynamic Range Grayscale to a Lower Resolution Display" par J.F. Boulter

On examine l'aptitude du filtrage spatial de type non linéaire à réduire la perte d'information quand un système d'affichage possède une plage de résolution dynamique inférieure à celle des niveaux de gris contenus dans l'image qu'on veut visionner. On obtient la compression de la plage dynamique par la réduction des basses fréquences et l'amplification des hautes fréquences en utilisant le filtrage homomorphe. Ce filtrage présuppose que l'image est multipliée par un arrière-plan à basse fréquence. Cette technique, décrite par Oppenheim pour les images formées par la réflexion de la lumière visible, est appliquée, par extension, aux images formées par l'émission de radiation IR et la transmission de rayons X et γ . A l'aide d'exemples, on illustre la compression de la plage dynamique sur ces trois types d'images en utilisant une filtration homomorphe et une procédure d'affichage sans paramètres libres à optimiser en fonction d'une image d'entrée. (NC)

DREV R-4074/77 (UNCLASSIFIED)

Bureau- Recherche et Développement, Ministère de la Défense nationale, Canada.
CRDV, C.P. 880, Courcellette, Qué. G0A 1R0

"Use of Spatial Filtering to Match Wide Dynamic Range Grayscale to a Lower Resolution Display" par J.F. Boulter

On examine l'aptitude du filtrage spatial de type non linéaire à réduire la perte d'information quand un système d'affichage possède une plage de résolution dynamique inférieure à celle des niveaux de gris contenus dans l'image qu'on veut visionner. On obtient la compression de la plage dynamique par la réduction des basses fréquences et l'amplification des hautes fréquences en utilisant le filtrage homomorphe. Ce filtrage présuppose que l'image est multipliée par un arrière-plan à basse fréquence. Cette technique, décrite par Oppenheim pour les images formées par la réflexion de la lumière visible, est appliquée, par extension, aux images formées par l'émission de radiation IR et la transmission de rayons X et γ . A l'aide d'exemples, on illustre la compression de la plage dynamique sur ces trois types d'images en utilisant une filtration homomorphe et une procédure d'affichage sans paramètres libres à optimiser en fonction d'une image d'entrée. (NC)

DREV R-4074/77 (UNCLASSIFIED)

Bureau- Recherche et Développement, Ministère de la Défense nationale, Canada.
CRDV, C.P. 880, Courcellette, Qué. G0A 1R0

"Use of Spatial Filtering to Match Wide Dynamic Range Grayscale to a Lower Resolution Display" par J.F. Boulter

On examine l'aptitude du filtrage spatial de type non linéaire à réduire la perte d'information quand un système d'affichage possède une plage de résolution dynamique inférieure à celle des niveaux de gris contenus dans l'image qu'on veut visionner. On obtient la compression de la plage dynamique par la réduction des basses fréquences et l'amplification des hautes fréquences en utilisant le filtrage homomorphe. Ce filtrage présuppose que l'image est multipliée par un arrière-plan à basse fréquence. Cette technique, décrite par Oppenheim pour les images formées par la réflexion de la lumière visible, est appliquée, par extension, aux images formées par l'émission de radiation IR et la transmission de rayons X et γ . A l'aide d'exemples, on illustre la compression de la plage dynamique sur ces trois types d'images en utilisant une filtration homomorphe et une procédure d'affichage sans paramètres libres à optimiser en fonction d'une image d'entrée. (NC)

DREV R-4074/77 (UNCLASSIFIED)

Research and Development Branch, Department of National Defence, Canada.
DREV, P.O. Box 880, Courcellette, Qué. G0A 1R0

"Use of Spatial Filtering to Match Wide Dynamic Range Grayscale to a Lower Resolution Display" by J.F. Boulter

We investigate the ability of nonlinear spatial filtering to reduce the amount of information that is lost when a display device has a lower dynamic range than the grayscale imagery to be displayed on it. Dynamic range compression can be achieved by attenuating an assumed low-frequency multiplicative background and enhancing higher frequencies using homomorphic filtering. This technique, originally described by Oppenheim for images formed by reflection of visible light, is extended to images formed by emission of IR radiation and transmission of X-ray and Y-ray radiation. Examples are given which show dynamic range compression of the three types of images achieved by means of a homomorphic filtering and display procedure having no free parameters which must be optimized to suit the particular input image. (U)

DREV R-4074/77 (UNCLASSIFIED)

Research and Development Branch, Department of National Defence, Canada.
DREV, P.O. Box 880, Courcellette, Qué. G0A 1R0

"Use of Spatial Filtering to Match Wide Dynamic Range Grayscale to a Lower Resolution Display" by J.F. Boulter

We investigate the ability of nonlinear spatial filtering to reduce the amount of information that is lost when a display device has a lower dynamic range than the grayscale imagery to be displayed on it. Dynamic range compression can be achieved by attenuating an assumed low-frequency multiplicative background and enhancing higher frequencies using homomorphic filtering. This technique, originally described by Oppenheim for images formed by reflection of visible light, is extended to images formed by emission of IR radiation and transmission of X-ray and Y-ray radiation. Examples are given which show dynamic range compression of the three types of images achieved by means of a homomorphic filtering and display procedure having no free parameters which must be optimized to suit the particular input image. (U)

DREV R-4074/77 (UNCLASSIFIED)

Research and Development Branch, Department of National Defence, Canada.
DREV, P.O. Box 880, Courcellette, Qué. G0A 1R0

"Use of Spatial Filtering to Match Wide Dynamic Range Grayscale to a Lower Resolution Display" by J.F. Boulter

We investigate the ability of nonlinear spatial filtering to reduce the amount of information that is lost when a display device has a lower dynamic range than the grayscale imagery to be displayed on it. Dynamic range compression can be achieved by attenuating an assumed low-frequency multiplicative background and enhancing higher frequencies using homomorphic filtering. This technique, originally described by Oppenheim for images formed by reflection of visible light, is extended to images formed by emission of IR radiation and transmission of X-ray and Y-ray radiation. Examples are given which show dynamic range compression of the three types of images achieved by means of a homomorphic filtering and display procedure having no free parameters which must be optimized to suit the particular input image. (U)

DREV R-4074/77 (UNCLASSIFIED)

Research and Development Branch, Department of National Defence, Canada.
DREV, P.O. Box 880, Courcellette, Qué. G0A 1R0

"Use of Spatial Filtering to Match Wide Dynamic Range Grayscale to a Lower Resolution Display" by J.F. Boulter

We investigate the ability of nonlinear spatial filtering to reduce the amount of information that is lost when a display device has a lower dynamic range than the grayscale imagery to be displayed on it. Dynamic range compression can be achieved by attenuating an assumed low-frequency multiplicative background and enhancing higher frequencies using homomorphic filtering. This technique, originally described by Oppenheim for images formed by reflection of visible light, is extended to images formed by emission of IR radiation and transmission of X-ray and Y-ray radiation. Examples are given which show dynamic range compression of the three types of images achieved by means of a homomorphic filtering and display procedure having no free parameters which must be optimized to suit the particular input image. (U)